
Multi-purpose serial data link

The MSDL (Multi-purpose Serial Data Link) I/O card introduced in Release 18 has greatly enhanced the I/O capability of the Meridian 1. Prior to the introduction of the MSDL card, a Meridian 1 system could support a total of 16 I/O ports. Now, a system can support up to 16 MSDL cards, each of which can be flexibly configured to support combinations of SDI, AML, CSL, and DCH on 4 ports, for a total of 64 I/O ports.

With this added flexibility and capability, however, comes additional complexity. Applications that were once relatively isolated from one another are now competing for resources on the MSDL, resulting in new interactions. Problems with a single application have the potential of disabling the entire card, jeopardizing other applications. Meanwhile, the MSDL, which was designed in X11 release 18, is being asked to do more as the core processor has become faster, with the introduction of the 68040 and 68060 processors. It becomes essential, then, to engineer the MSDL card to minimize the risk of performance problems. This document provides guidelines to help the user in engineering the MSDL.

Note that these engineering guidelines assume normal traffic consisting of valid call processing and administrative messages. For example, engineering rules cannot prevent a piece of equipment on the network from malfunctioning and generating spurious messages which overload the MSDL. At this point the recovery mechanism becomes essential. The mechanism should be graceful, not requiring manual intervention, and should provide as much diagnostic information as possible, to help isolate the root cause of the problem. Refinements and improvements to the recovery mechanisms have been introduced over various software releases.

Overview

Engineering the MSDL requires an understanding of the end-to-end performance characteristics of the system, including the Meridian 1, MSDL, link, and terminating or originating node. Outgoing messages originate from the Meridian 1 CPU, are passed to the MSDL, and travel across the appropriate link to the destination. In equilibrium, or over a relatively long period of time, i.e. on the order of several minutes, the Meridian 1 cannot generate messages faster than the MSDL processor can process them, than the link can transmit them, or than the destination can process them. Otherwise, messages will build up at the bottleneck and will eventually be lost. The entity with the lowest capacity will be the system bottleneck. For very short periods of time, however, one or more entities is able to send messages at a higher rate than the system bottleneck, since buffers are available to queue the excess messages. These periods are referred to as bursts. The length of the burst and the size of the burst that can be supported depend on the sizes of the buffers.

Thus, to properly engineer a system, two areas are considered: equilibrium or steady-state performance which requires an analysis of the CPU processing capacity of the various components of the system along with link bandwidth; and burst performance which requires an analysis of the buffer utilization of the system. The equilibrium analysis assumes 30% peakedness which is consistent with models for the Meridian 1 CPU.

The applications which will be discussed here are: DCH, CSL/AML, and SDI. The Meridian 1 CPUs considered include the 68030, 68040, and 68060 processors.

Section provides a brief overview of the MSDL architecture. [“DCH” on page 176](#) through [“SDI” on page 192](#) describe general conditions for equilibrium and peak engineering for key applications. A step-by-step procedure for engineering the MSDL is provided in [“MSDL Engineering Procedure” on page 197](#). Several examples of the engineering procedure are given in [“Examples” on page 205](#).

MSDL Architecture

The MSDL processor is a 68020 processor. The MSDL and Meridian 1 exchange messages using an SRAM and interrupt scheme. To prevent any one application from tying up buffer resources, a flow control mechanism is defined at the SL-1 and MSDL/MISP interface level, where SL-1 denotes the call processing software running on the Meridian 1 core CPU. The flow control mechanism is based on the common window mechanism in which the number of messages outstanding in the transmit or receive direction per socket, or port, cannot exceed $T(K)$ or $R(K)$, respectively. In the transmit direction, for example, a message is considered outstanding from the time the SL-1 software writes it into the transmit ring until all processing of the message by the MSDL is completed. Currently $T(K)$ and $R(K)$ are both set at 30. Each application must queue messages if the flow control threshold is exceeded. Typically, the Meridian 1 task also has a buffer for messages.

An overload control threshold is also implemented in the incoming direction to protect the Meridian 1 CPU from excess messages. Prior to Release 23, this threshold was set at 100 messages per second. MSDL305 was printed if 100 messages in 2 seconds was exceeded, and MSDL306 was printed and the card disabled if the 100 messages in 1 second threshold was exceeded. With Release 23, to account for the new, faster processors, the thresholds have been changed so that MSDL304 is printed if 100 messages in 2 seconds is exceeded, MSDL305 is printed if 200 messages in 2 seconds is exceeded, and MSDL306 is printed and the card is disabled if 300 messages in 2 seconds is exceeded. In both cases Background Audit will bring the MSDL back up if no problems are found.

Several software tasks exist on the MSDL. Layer 1 message processing operates at the highest priority. If the link is noisy, Layer 1 processing may starve the Layer 2 and Layer 3 processing tasks, resulting in buffer overflows. If such a problem is suspected, the Protocol Log (PLOG) should be examined. PLOG reporting is requested in Overlay 96, as described in the *X11 input/output guide*.

DCH

For pre-Release 20 primary rate interfaces including Custom DMS-100, Custom 5ESS, 4ESS, DMS-250, and Meridian 1 to Meridian 1, the MSDL software performs Layer 1 and Layer 2 processing. These interfaces are referred to as pre-R20 interfaces and are the interfaces supported on the DCHI card. Based on real time measurements on an MSDL, the rated capacity of the MSDL is 87 D-channel messages per second for pre-R20 interfaces. For interfaces developed in Release 20 and later, including NI-2, Q-SIG, and Euro-ISDN, Layer 3 processing is also performed on the MSDL, so the MSDL performs some functions previously performed by the Meridian 1 core processor, thus reducing the capacity on the MSDL. These interfaces will be referred to as R20+ interfaces. The steady state message rate allowable for D-channel messages is 29 msg/sec for R20+ interfaces.

The SL-1 software output queue for DCH messages is the Output Buffer (OTBF) which is user configurable for between 1 and 127 buffers in Overlay 17. This is a single system resource which is shared by all D-channels.

Beginning with Release 22, it is possible to define overload thresholds for R20+ interfaces on a per-D-channel basis. The ISDN_MCNT (ISDN message count), defined in Overlay 17, specifies the number of ISDN Layer 3 call control messages allowed per 5-second interval. Overload control thresholds can be set on a per D-channel basis, ranging from 60 to 350 messages in a 5 second window, with a default of 300 messages. If the overload control threshold is exceeded, DCH421 is output. When the message rate exceeds the threshold for two consecutive 5 second periods, overload control is invoked and new incoming call requests are rejected by the Layer 3 protocol control in the third 5 second time interval. Layer 3 will resume accepting new calls at the end of the third time interval. This flexibility allows the user to regulate the MSDL processing required by a specific R20+ DCH port. Note that the default value implies no overload control since 300 messages/5 seconds exceeds the rated capacity of 29 messages/second.

PRI Network

Equilibrium Analysis

A D-channel can be configured to support up to 383 B-channels (or 382 with a back-up D channel) on a T1 or 480 B-channels on an E1. The bandwidth available for messages is 64 kbps. Assumptions for a typical application are: 8 messages/call, 29 bytes/message, including 18 bytes of Layer 3 data and 11 bytes of Layer 2 overhead, 28 centi-call seconds (CCS)/trunk, and 180 second Average Hold Time (AHT)/call. The Meridian 1 capacity is derived from its call carrying capacity for 100% incoming PRI calls.

Under the traffic assumptions described above, the MSDL is able to support basic call processing messages for 4 D-channels under normal operation (see [Table 60](#)).

Table 60
Steady-state requirements and capacities per D-channel (outgoing and incoming)

	Requirement msg/sec	Meridian 1 CP capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec	
68030 CPU	13(T1)/16(E1)	77	87	212 input 212 output	Limited by traffic requirements
68040 CPU	13(T1)/16(E1)	121	87	212 input 212 output	Limited by traffic requirements
68060 CPU	13(T1)/16(E1)	161	87	212 input 212 output	Limited by traffic requirements

Peak Analysis

When there is a link re-start, STATUS messages are sent to all trunks with established calls. Since the SL-1 software task does not implement flow control on this mechanism, a burst of up to several hundred messages can be sent to the MSDL, exceeding MSDL flow control thresholds. When this happens, messages back up on the OTBF buffer, possibly resulting in buffer overflow, as indicated by DCH1030 messages. OTBF overflow is also possible after an initialization since a burst of messages is sent to each D-channel in the system, and the OTBF is a shared system resource.

The Meridian 1 capacity is significantly higher in this scenario than in the previous one because it is sending out D-channel messages which do not involve call processing. MSDL and Link capacities are also higher because, for equilibrium analysis, some capacity is reserved for peaking.

Table 61 illustrates the worst case scenario for a single D-channel. If the Meridian 1 sends messages at its peak rate, OTBF buffer overflow is possible. Also, once the messages are sent, a burst of responses can be expected in the incoming direction, resulting in additional congestion at the MSDL.

Table 61
Peak requirements per D-channel (outgoing)

	Burst Size	Meridian 1 capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec	
68030 CPU	382(T1)/480(E1)	195	113	276 output	MSDL is bottleneck
68040 CPU	382(T1)/480(E1)	308	113	276 output	MSDL is bottleneck
68060 CPU	382(T1)/480(E1)	410	113	276 output	MSDL is bottleneck

This situation also occurs when a back-up D-channel becomes active, since STATUS messages are exchanged to resynchronize the link.

To reduce the possibility of this problem occurring, limit the number of B-channels supported by a D-channel, separate D-channels onto several MSDL cards so that message bursts are not being sent to four ports on the same MSDL after initializations, and increase the size of OTBF to the maximum value of 127.

B-Channel Overload

In an ACD environment in which the number of ACD agents plus the maximum ACD queue length is considerably less than the number of B-channels available for incoming calls, a burst of incoming messages may impact the performance of the MSDL as well as the Meridian 1 via the following mechanism: Calls from the CO terminate on a specified ACD queue. When the destination is busy, i.e. the destination set is busy or the ACD queue has reached its maximum limit of calls, the Meridian 1 immediately releases the call. The CO will immediately present another call to the same destination, which is released immediately by the PBX, etc.

In release 23, the B-Channel Overload Control feature is introduced to address this problem by delaying the release of an ISDN PRI call by a user-configurable time when the call encounters a busy condition. The delay in releasing the seized B-channel prevents a new call from being presented on the same B-channel, decreasing the incoming call rate. The timer BCOT is set in Overlay 16, and falls in the range 0 to 4000 msec.

ISL Network

In an ISL application, a modem is used to transmit ISDN signaling messages. Baud rates are user configurable at the standard RS232/RS422 rates: 300, 1200, 2400, 4800, 9600, and 19200 bps (see [Table 62](#)). In this case, the modem baud rate constraint can be the limiting constraint. The messages/second that can be supported by the baud rates are given below, where the values allow for 30% peakedness.

The B-channels that can be supported assume the messaging required for a typical application as described in [“Equilibrium Analysis” on page 177](#).

Table 62
ISL link capacities

Modem Baud Rate	Link capacity msgs/sec	B-channels that can be supported
300	1 input 1 output	46
1200	4 input 4 output	180
2400	7 input 7 output	316
4800	15 input 15 output	382(T1)/480(E1)
9600	29 input 29 output	382(T1)/480(E1)
19200	58 input 58 output	382(T1)/480(E1)

For the baud rates listed in Table 62, the link will be the limiting constraint. The potential peak traffic problems described in Section apply here as well, to an even greater extent since the rate mismatch between the Meridian 1 and the system bottleneck, now the link instead of the MSDL, is greater. To minimize the risk, set the baud rate as high as possible.

VNS Network

The discussion concerning ISL networks applies to VNS networks as well. Prior to release 22, up to 100 VDNs was supported. With release 22 this number was expanded to 4000.

NACD Network

A Network ACD (NACD) network is difficult to engineer since performance depends on specific network configuration details including connectivity, routing tables, the number of nodes, the number of queues at each node, and calling patterns.

Diverting calls in NACD is controlled by Routing Tables with timers. Calls diverted by NACD can be answered by the Source ACD DN or any one of up to 20 Target ACD DNs. Each Target can have an individual timer defined, from 0 to 1800 seconds. By using ISDN D-channel messaging to queue Call Requests at remote Target ACD DNs, voice calls are not physically diverted until an idle agent is reserved for that call at the remote Target node.

It is recommended that the Routing Table be designed so that Call Requests cascade to the network with the timers staggered. The node that is most likely to have available agents should have the smallest timer value. Otherwise Call Requests will flood the network, resulting in inefficient use of network and real time resources.

An Active Target is available to accept NACD calls, while a Closed Target is closed to incoming calls. When calls in the Call Request queue exceed the Call Request Queue Size (CRQS) threshold, the status changes to Closed. A Status Exchange message is sent from the Target node to the Source ACD DNs indicating the new status. The Target ACD DN remains Closed to further network call requests until the number of calls in the queue is reduced by the Flow Control Threshold (FCTH).

Equilibrium Analysis

At the source node, for each call queued to the network but not answered, 4 messages are exchanged. For each call queued to the network and answered, 11 messages are exchanged. Likewise, at the target node, a network call that is queued but not answered requires 4 messages while a call that is queued and answered requires 11 messages. Messages average 31 bytes.

From a single D-channel perspective, the most difficult network topology is a star network in which each agent node is connected to a tandem node (see [Table 63](#)). All messages to the other nodes are sent across the D-channel connected to the tandem node. As an example, consider a site with 2000 calls arriving locally during the busy hour. The timers in the Routing Table are staggered so that 1000 are answered locally without being queued to the network, 500 are answered locally after being queued to an average of two network target queues, and 500 are answered in the network after being queued to an average of four network target queues. Meanwhile, 200 Logical Call Requests arrive from the network, of which 100 calls are answered.

Table 63
Steady-state requirements and capacities per D-channel with staggered timers
(outgoing and incoming)

	Requirement msg/sec	Meridian 1 CP capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec	
68030 CPU	5	77	87	212 input 212 output	Limited by traffic requirements
68040 CPU	5	121	87	212 input 212 output	Limited by traffic requirements
68060 CPU	5	161	87	212 input 212 output	Limited by traffic requirements

For this same network, assume now that the timers in the Routing Table are not staggered; instead, Logical Call Requests are broadcast to the four target nodes in the network as soon as calls arrive at the local node. Also assume that a total of 4000 calls arrive elsewhere in the network, and are queued at local ACD DNs. Even if the calls are answered exactly where they were before, the number of messages exchanged will increase significantly, to the values provided in Table 64, using the following calculations:

- 1500 calls queued on 4 ACD DNs and not answered * 4
msgs/call/DN = 24000 msgs
- 500 calls answered * 11 msgs/call = 5500 msgs
- 500 calls queued on 3 ACD DNs and not answered * 4
msgs/call/DN=6000 msgs
- 3900 network calls queued on local DN and not answered * 4
msgs/call=15600 msgs
- 100 network calls answered * 11 msgs/call=1100 msgs
- Total 52200 msgs/hr
- $(52200 \text{ msgs/hr}) / (3600 \text{ secs/hr}) = 14.5 \text{ msgs/sec}$

Table 64
Steady-state requirements and capacities per D-channel with immediate broadcast of Logical Call Requests (outgoing and incoming)

	Requirement msg/sec	Meridian 1 CP capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec	
68030 CPU	14.5	77	87	212 input 212 output	Limited by traffic requirements
68040 CPU	14.5	121	87	212 input 212 output	Limited by traffic requirements
68060 CPU	14.5	161	87	212 input 212 output	Limited by traffic requirements

Peak Analysis

When the CRQS threshold is reached, the target queue will broadcast messages to the source ACD DN's informing them that it will no longer accept calls. The size of this outgoing burst of messages depends on the number of source ACD DN's in the network.

Once the FCTH threshold is reached, another Status Exchange message is sent. At that point, Logical Call Request messages are sent by the Source ACD DN's. While the target queue has been closed, many calls may have queued at source ACD DN's, resulting in a burst of Logical Call Request messages once the DN becomes available.

Unlike the PRI network case, there is no specific worst case scenario for peakedness. The examples in [Tables 65](#) and 8 are based on a 5 node network, where each node has three source ACD DN's.

Table 65
Peak requirements for NACD messages (outgoing)

	Burst Size	Meridian 1 capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec	
68030 CPU	12	195	113	258 output	MSDL is bottleneck
68040 CPU	12	308	113	258 output	MSDL is bottleneck
68060 CPU	12	410	113	258 output	MSDL is bottleneck

Table 66
Peak requirements for NACD messages (incoming)

	Burst Size	Meridian 1 capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec	
68030 CPU	40	195	113	258 input	MSDL is bottleneck
68040 CPU	40	308	113	258 input	MSDL is bottleneck
68060 CPU	40	410	113	258 input	MSDL is bottleneck

If CRQS values are set high, many messages will be exchanged, with the network emulating a single virtual queue. If the CRQS values are lowered, fewer Call Requests will be sent across the network, however, average source delays may be increased. If FCTH levels are set too low, target nodes can ping pong between Active and Closed states, resulting in network congestion and excessive real time utilization. However, if FCTH levels are set too high, a target node may be inundated with Logical Call Request messages once it becomes available. CRQS is configurable for the range [0, 255], while FCTH is configurable for the range [10, 100]. Since the impact of these parameters is so configuration dependent, it is beyond the scope of this document to make recommendations on how to configure them. They should be determined as part of the custom network design process. Contact your local Northern Telecom representative for network engineering services.

Impact of Proper Engineering of B-channels

In the NACD environment another problem arises when insufficient B-channels are configured across the network. When an agent becomes available, an Agent Free Notification message is sent to the source node. An ISDN Call Setup message is sent from the source node to the target node. Since no B-channel is available, the agent reservation timer expires, and an ISDN Cancellation Message is sent from the target node to the source node and an ISDN Cancellation Acknowledge message is sent from the source node to the target node. At this point, the agent is still free, so the process repeats until a trunk becomes available or the target closes. This scenario results in a significant amount of message passing.

Parameter Settings

The following are parameters that can be configured in Overlay 17 for Meridian 1 D-channels. They are listed with their input range and default value in ().

OTBF 1 - (32) - 127: Size of output buffer for DCH

This parameter configures how many output buffers are allocated for DCH messages outgoing from the Meridian 1 CP to the MSDL card. The more that are created, the deeper the buffering. Normally a message created in a buffer is sent to the MMIH (Meridian MSDL Interface Handler) and copied into the ring. If the ring is flow controlled, the message occupies a buffer until it can be sent. For systems with extensive D-channel messaging, such as call centers using NACD, the parameter should be set at 127. For other systems with moderate levels of D-channel messaging, OTBF should be set at the smaller of the following two quantities: Total B-channels - (30 * MSDL cards with D-channels) or 127.

For example, if a system in a standard office environment is configured with 7 T1 spans, 2 D-channels which are located on two different MSDLs, and 2 back-up D-channels, the total number of B-channels is $(7*24)-4=164$. OTBF should be configured to be the smaller of $164-(30*2)=104$ and 127 which is 104.

T200 2 - (3) - 40: Maximum time for acknowledgment of frame (units of 0.5 secs)

This timer defines how long the MSDL's Layer 2 LAPD will wait before it retransmits a frame. If it doesn't receive an acknowledgment from the far end for a given frame before this timer expires, it will retransmit a frame. Setting this value too low can cause unnecessary retransmissions. The default of 1.5 seconds is long enough for most land connections. Special connections, over radio, for instance, may require higher values.

T203 2 - (10) - 40: Link Idle Timer (units of seconds)

This timer defines how long the Layer 2 LAPD will wait without receiving any frames from the far end. If no frames are received for a period of T203 seconds, the Layer 2 will send a frame to the other side to check that the far end is still alive. The expiration of this timer causes the periodic “RR” or Receiver Ready to be sent across an idle link. Setting this value too low causes unnecessary traffic on an idle link. However, setting the value too high will delay the system from detecting that the far end has dropped the link and initiating the recovery process. The value should be higher than T200. It should also be coordinated with the far end so that one end does not use a small value while the other end uses a large value.

N200 1 - (3) - 8: Maximum Number of Retransmissions

This value defines how many times the Layer 2 will resend a frame if it doesn't receive an acknowledgment from the far end. Every time a frame is sent by Layer 2, it expects to receive an acknowledgment. If it does not receive the acknowledgment, it will retransmit the frame N200 times before attempting link recovery action. The default (3) is a standard number of retransmissions and is enough for a good link to accommodate occasional noise on the link. If the link is bad, increasing N200 may keep the D-channel up longer, but in general this is not recommended.

N201 4 - (260): Maximum Number of Octets (bytes) in the Information Field

This value defines the maximum I-frame (Info frame) size. There is no reason to reduce the number from the default value unless the Meridian 1 is connected to a system that does not support the 260-byte I-frame.

K 1 - (7): Maximum number of outstanding frames

This value defines the window size used by the Layer 2 state machine. The default value of 7 means that the Layer 2 state machine will send up to 7 frames out to the link before it stops and requires an acknowledgment for at least one of the frames. A larger window allows for more efficient transmission. Ideally, the Layer 2 will receive an acknowledgment for a message before reaching the K value so that it can send a constant stream of messages. The disadvantage of a large K value is that more frames must be retransmitted if an acknowledgment is not received. The default value of 7 should be sufficient for all applications. The K value must be the same for both sides of the link.

ISDN_MCNT (ISDN Message Count) 60 - (300) - 350: Layer 3 call control messages per 5 second interval

Beginning with release 22, it is possible to define overload thresholds for R20+ interfaces on a per-D-channel basis. This flexibility allows the user to regulate the MSDL processing required by a specific R20+ DCH port. The default value of 300 messages/5 seconds is equivalent to allowing a single port to utilize the full real time capacity of an MSDL. To limit the real time utilization of a single R20+ DCH port to $1/n$ of the real time capacity of the MSDL, for $n > 1$, set ISDN_MCNT to $(300 / n) * 1.2$ where the 1.2 factor accounts for the fact that peak periods on different ports are unlikely to occur simultaneously. For example, to limit a single port to $1/3$ of the processing capacity of the MSDL, ISDN_MCNT is set to $(300 / 3) * 1.2 = 120$.

If the ISDN_MCNT threshold is exceeded for one 5 second period, error message DCH421 is printed. If the threshold is exceeded for two consecutive periods, incoming call requests arriving in the third 5 second interval are rejected by the MSDL Layer 3 software. At the end of the third 5 second interval, Layer 3 will resume accepting incoming call requests.

AML

The Application Module Link (AML) provides the connection between the Meridian 1 and the CCR, Meridian Link, or Meridian 911 module. The current maximum speed for the link is 19200 baud. CCR is the application addressed here because it is the one that results in the highest level of messaging. The amount of messaging involved depends on the complexity of call handling. Simple call handling results in approximately 10 messages per call, with an average of 45 bytes/message. Statistics messages are sent from the Meridian 1 to the CCR module every 4 seconds for ACD DN's referenced in the CCR variable table or scripts. Thus messaging levels depend not only on the number of calls handled but on the number of ACD DN's with statistics configured. **Current recommendations are that a system be limited to 80 ACD DN's with statistics.**

On the Meridian 1, messages queue in the CSQI and CSQO buffers, command status queue input and output buffers, which are configurable in Overlay 17.

Equilibrium Analysis

For equilibrium analysis, we focus on calls, and assume ten ACD DN's sending statistics messages. The Meridian 1 capacity assumes an inbound call center with simple CCR treatment on 100% of the calls, and Meridian MAX.

For large systems, the CCR module capacity is the system bottleneck see [Table 67](#)). Since there is no flow control or overload control available to protect the CCR module, it is essential that the system be engineered to ensure that the CCR module is not overloaded. Otherwise, link failures or other CCR performance problems may result. To engineer the CCR module, refer to the *Meridian Link/Customer Controlled Routing Engineering Guide 553-3211-520*.

Table 67
Steady-state requirements and capacities per AML (outgoing and incoming)

	Meridian 1 CP capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec	CCR capacity msg/sec (167 module)	
68030 CPU	39	107	41 input 41 output	46	68030 CPU bottleneck
68040 CPU	56	107	41 input 41 output	46	CCR bottleneck
68060 CPU	74	107	41 input 41 output	46	CCR bottleneck

Peak Analysis

Since message bursts are most likely to cause buffer overflow, we consider the system with 80 ACD DN's sending statistics messages every 4 seconds. Recall that this is the maximum recommended number for ACD DN's sending statistics. The Meridian 1 capacity is based on the real time required to process CCR statistics messages (see [Table 68](#)).

Table 68
Peak capacities for CCR statistics messages per AML (outgoing)

	Burst Size	Meridian 1 capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec	CCR capacity msg/sec (167 module)	
68030 CPU	80	420	139	53	60	AML bottleneck
68040 CPU	80	692	139	53	60	AML bottleneck
68060 CPU	80	920	139	53	60	AML bottleneck

In this scenario, the AML link is the bottleneck. Messages will begin to queue in the MSDL output buffers and possibly the CSQO buffers, if there are many ACD DN's sending statistics messages.

The AML link will disable if 10 consecutive messages do not receive a response within a 4 second window. The CSA105 message is normally output when this occurs. If a message arrives immediately after the statistics messages for the 80 ACD DN's are generated, it may be queued behind these 80 statistics messages. For 80 messages, processing time at the MSDL, queueing time for the AML, and processing time at the CCR module add up to approximately 3 seconds, so it is easy to understand how the 4 second threshold might be exceeded if the MSDL is also processing messages from other applications.

In Meridian Link applications, similar types of problems may occur when the host is too slow and becomes the system bottleneck.

Parameters

On the Meridian 1 side, AML messages are queued in the CSQI/CSQO buffers, which are shared with the CSL. The maximum configurable size of each is 25% of the number of call registers in the system or 255. It is recommended that 68040 CPU and 68060 CPU systems configure the CSQI and CSQO buffers to be 255, while 68030 CPU systems configure 200. CSQO and CSQI sizes are configured in Overlay 17.

The AML flow control parameters MCNT and INTL are also set in Overlay 17. This flow control mechanism limits the number of messages sent from the CCR to the Meridian 1 to MCNT [5,9999] in the time interval INTL [1,12] where INTL is measured in units of 5 seconds. When this threshold is violated for one interval, a warning message is sent to CCR requesting that it slow down. If the threshold is violated for two consecutive periods, CCR rejects all new calls back to the Meridian 1 where they will receive default treatment. No new calls will be accepted until the level of traffic is reduced to an acceptable level. If the threshold is exceeded for three consecutive periods, all inbound traffic will be lost. If inbound traffic continues, the link will fail.

Recommended settings for MCNT and INTL are listed in [Table 69](#).

Table 69
Recommended AML flow control values

	MCNT	INTL
68030 CPU	600	1
68040 CPU	900	1
68060 CPU	1400	1

This mechanism was originally designed to protect the Meridian 1 from overload. With the faster processors, this flow control threshold is unlikely to be exceeded.

SDI

An asynchronous serial data interface was provided on the MSDL card starting with release 19. Capabilities include interface to TTYs, printers, modems, and CRTs, High Speed Link (HSL) for ACD, Auxiliary Processor Link (APL) for ACD, ACD-C package displays and reports, and CDR TTYs. An SDI port is only configurable on Port 0 of an MSDL. Therefore, only one SDI port can be configured on an MSDL.

Normally, in the output direction, the SDI Application will pass any character received from the Meridian 1 to the Layer 1 Driver to be sent out over the interface. If XON/XOFF Handling is enabled for printing, the SDI Application will buffer up to 500 characters once an XOFF is received. The Meridian 1 is not aware that an XOFF has been received. After the buffer is full, if further output is received, the oldest data will be discarded. Output resumes when an XON is received or 1 minute has passed since the output was halted by an XOFF. At this point, the contents in the buffer will be emptied first, followed by output from the Meridian 1. If any data has been discarded, an error message will be sent.

In the input direction, every character received by the Layer 1 Driver will be passed to the SDI Application. The SDI Application will echo any input character unless it is told not to by the Meridian 1. In Line Editing Mode, the SDI Application will buffer a line of up to 80 characters which can be edited before being sent to the Meridian 1.

Under certain conditions, control characters can cause messages to ping pong between a modem or printer and the Meridian 1, resulting in MSDL305 or MSDL306 conditions. To avoid these situations, configure modems in dumb mode and disable printer flow control.

The Meridian 1 input buffer is the TTY input buffer which can store 512 characters. The Meridian 1 output buffer is the TTY output buffer which can store 2048 characters.

CDR

CDR records are available in two formats: *FCDR=old* and *FCDR=new*. A typical record for the old format is 100 bytes long while a typical record for the new format is 213 bytes long (see [Table 70](#)). Due to the nature of the SDI interface, characters are output one at a time, resulting in 100 messages and 213 messages generated for *FCDR=old* and *FCDR=new*, respectively. Each message requires 10 bits. Based on real time measurements, the MSDL rated capacity for processing CDR messages is 16,631 messages/second.

Table 70
Link capacities for CDR application (outgoing)

Modem Baud Rate	Link capacity msg/sec (peak)	Calls/Hour for <i>FCDR=old</i>	Call/Hour for <i>FCDR=new</i>
300	30	831	390
1200	120	3323	1560
2400	240	6646	3120
4800	480	13292	6241
9600	960	26585	12481
19200	1920	53169	24962
38400	3840	106338	49924

Equilibrium Analysis

The Meridian 1 capacity for messages per second is conservatively based on the assumption of 100% outgoing calls with *FCDR=new*. Typically, CDR records are not generated for 100% of the calls (see [Table 71](#)).

Table 71
Steady state requirements for CDR application (outgoing)

	Meridian 1 CP capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec	
68030 CPU	975	16631	See Table 70	9600 baud recommended
68040 CPU	1537	16631	See Table 70	19200 baud recommended
68060 CPU	2044	16631	See Table 70	19200 baud recommended

Peak Analysis

Since each character is sent as a separate message, every time a CDR record is sent, a traffic peak is generated. In [Table 72](#) we consider *FCDR=new*.

Table 72
Peak requirements for *FCDR=new* (outgoing)

	Burst Size	Meridian 1 capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec	
68030 CPU	213	1474	21620	See Table 70	19200 baud recommended
68040 CPU	213	2323	21620	See Table 70	38400 baud recommended
68060 CPU	213	3090	21620	See Table 70	38400 baud recommended

MSDL real time capacity is not the bottleneck in this case. However, to prevent system buffers from building up, the recommended baud rate should be set. If a lower baud rate is chosen, assume that the CDR application will frequently be in a state of flow control. Note that this is true even if the steady state message rate is low, due to the nature of the SDI interface.

The burst sizes will be even greater if CDR is configured with queue records for incoming ACD calls.

Meridian MAX

The Meridian 1 communicates with Meridian MAX via the HSL (High Speed Link) using 8 bits plus one stop bit. Prior to MAX 8, the HSL bandwidth was 9600 baud. With MAX 8, 19,200 baud is available. Unlike the CDR application, MAX reports are not sent out character by character. The MAX report for a simple call is 5 messages of 20 bytes. The Meridian MAX module capacities are given in [Table 73](#).

Table 73
Capacity of MAX Module in simple calls

	MAX Capacity (simple calls/hour)	MAX Capacity (msg/sec)
Meridian MAX Rls 4	10000	14
Meridian MAX Rls 5	15000	21
Meridian MAX Rls 6	15000	21
Meridian MAX Rls 7	15000	21
Meridian MAX Rls 8	30000	42

Equilibrium Analysis

The Meridian 1 capacity requirements are derived assuming a simple inbound call center with all calls answered by ACD agents and MAX reporting on all calls (see [Table 74](#)). Incoming CDR is not turned on.

Table 74
Steady-state requirements for Meridian MAX (outgoing)

	Meridian 1 CP capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec (9600/19200)	MAX module capacity
68030 CPU	23	87	41/82	See Table 73
68040 CPU	36	87	41/82	See Table 73
68060 CPU	48	87	41/82	See Table 73

The 19,200 baud option for HSL is recommended for 68060 CPU systems.

Note that Meridian MAX messages queue on call registers. There is no software limit to the number of call registers that can be used to store MAX reports. If the PBX becomes overloaded, MAX messages may build up in the call registers, potentially impacting call processing since call registers may not be available for call setup.

Peak Analysis

Complex calls may require 15 or more messages per call. Depending on the configuration, the MSDL or the link could be the bottleneck. In either of these cases, messages queue in the system buffers (see [Table 75](#)).

Table 75
Peak requirements for MAX (outgoing)

	Burst Size	Meridian 1 capacity msg/sec	MSDL capacity msg/sec	Link capacity msg/sec (9600/19200)	MAX capacity
68030 CPU	15	92	87	53/106	See Table 73
68040 CPU	15	142	87	53/106	See Table 73
68060 CPU	15	189	87	53/106	See Table 73

MSDL Engineering Procedure

It is important to engineer MSDLs in the context of engineering the entire system, as discussed in previous sections. Refer to *Meridian 1 capacity engineering 553-3001-149* and *Traffic measurement: Formats and output 553-2001-450* for information on real time engineering of the Meridian 1. In all cases with a user configurable link rate, it is essential that the link be configured so that the rate is high enough to support steady state requirements and some peakedness. Otherwise these applications messages will occupy system buffers, increasing the chance of buffer overflow.

Table 76 is the high-level worksheet for analysis of MSDL capacity. The appropriate values can be derived from Table 77 through Table 82.

Table 76
MSDL Engineering Worksheet

Port	Application	Real Time Required	Peak Buffer Usage Outgoing	Peak Buffer Usage Incoming
0	_____	_____	_____	_____
1	_____	_____	_____	_____
2	_____	_____	_____	_____
3	_____	_____	_____	_____
Total		_____	_____	_____

Assuming 30% peakedness for the applications, the total real time required should be less than 2,770,000 msec. The projected real time utilization of the MSDL is given by

$$\text{MSDL_RTU} = \text{Total Real Time Required} / 2,770,000.$$

It is recommended that peak buffer usage be less than 60 in each direction. As the peak buffer usage increases over 60, the likelihood of an intermittent buffer full problem increases.

Application Engineering

The following sections provide procedures for calculating the real time required on the MSDL for various applications. In any of these cases, if the calls/hour value is known, insert that value into Column A. Otherwise, follow the guidelines provided. Values in parentheses () are default values. For example, the default number of calls/hr/trunk is 15.6. The value in Column E should be inserted in the Real Time Required column of Table 76 and the appropriate Peak Buffer Usage values should be inserted in the corresponding Peak Buffer Usage columns of Table 76.

DCH Applications

If several applications share a D-channel, the final real time requirements for the applications should be added and then entered in the appropriate entry in Table 77.

Table 77
MSDL real time requirements for D-channel applications

DCH	calls/hr A	msgs/call B	msgs/hr C=A*B	msec/msg** D	msec E=C*D
ISDN Network	trunks/DCH (see note)* calls/hr/trunk (15.6)= _____	8	_____	pre-R20: 8.8 R20+: 26.5	_____
NACD	NACD agents (see note)* calls/hr/agent (18.3)= _____	30	_____	pre-R20: 8.8	_____
NMS	NMS ports (see note)* calls/hr/port (65)= _____	10	_____	pre_R20: 8.8	_____
Note: For clarification of the terms “pre-R20” and “R20+,” refer to “DCH” on page 176					

The calculations described for NACD provide a simplified approximation of a “typical” NACD network. If call flows can be predicted or estimated, they can be used to develop a more accurate model using the number of messages described in Section . When this is done, the msgs/hr is computed directly, so columns A and B are not used. See [“Examples” on page 205](#) for a detailed example of how this can be done.

If a live system is being modeled, add the “number of all incoming messages received on the D-channel” and the “number of all outgoing messages sent on the D-channel” field from a busy hour TFS009 report to derive the entry for Column C. See *Traffic measurement: Formats and output 553-2001-450* for details.

Table 78
MSDL peak buffer requirements for D-channel applications

DCH	Outgoing	Incoming
ISDN Network	B-channels/DCH= _____	B-channels/DCH= _____
NACD	Source ACD DNs + 5 = _____	Network congestion level: Low: 10 Medium: 20 High: 30
NMS	10	10

In the case of an ISL D-channel, ensure that the baud rate of the connection is greater than

$$(C \text{ msgs/hr} * 29 \text{ bytes/msg} * 8 \text{ bits/byte}) / 3600 \text{ sec/hr}$$

where C comes from column C in Table 77.

If the baud rate is too low to meet requirements, performance of the entire MSDL card may be jeopardized since 30 of the MSDL output buffers will be occupied with ISL D-channel messages and the real time spent processing these messages will increase due to additional flow control and queueing logic.

Depending on the application, it may be too conservative to engineer an MSDL for link restarts. In that case, the ISDN Network peak outgoing and incoming buffer requirements can be set at 5 for 68030 CPU, 10 for 68040 CPU, and 15 for 68060 CPU systems.

AML Applications

If an existing system is being modeled, add the number of incoming messages, messages in the IMMSG category, and outgoing messages, messages in the OMSG category, from a busy hour TFS008 report and enter the value in Column C. For a quick approximation of the number of incoming messages, add the number of messages of priority 1 to 4, as provided in TFS008. For more details, refer to *Traffic measurement: Formats and output 553-2001-450*.

Table 79
MSDL real time requirements for AML applications

AML	calls/hr A	msgs/call B	msgs/hr C=A*B	msec/msg D	msec E=C*D
CCR	agents * calls/agent/hr (18.3)* %calls with CCR=_____	simple: 10 medium: 20 complex: 30	A*B + 900 ACD DNs w/ statistics = _____	7.2	_____
HER/AST	agents * calls/agent/hr (18.3)* % calls with HER/AST=_____	10	_____	7.2	_____
M911	M911 agents * calls/agent/hr (18.3)= _____	6	_____	7.2	_____
Meridian Mail voice mail	MM ports * calls/hr/port (65)= _____	10	_____	7.2	_____
Meridian Mail voice menu	agents * calls/agent/hr (120) = _____	10	_____	7.2	_____
Meridian Mail announcements	agents * calls/agent/hr (150)= _____	5	_____	7.2	_____

Table 80
MSDL peak buffer requirements for AML applications

AML	Outgoing	Incoming	Minimum Baud Rate
CCR	CDNs with statistics= _____	68030 CPU: 10 68040 CPU: 15 68060 CPU: 20	(msgs/hr * 45 bytes/msg * 8 bits/byte)/(3600 sec/hr)=_____
HER/AST	68030 CPU: 5 68040 CPU: 8 68060 CPU: 12	68030 CPU: 5 68040 CPU: 8 68060 CPU: 12	(msgs/hr * 45 bytes/msg * 8 bits/byte)/(3600 sec/hr)=_____
M911	68030 CPU: 2 68040 CPU: 3 68060 CPU: 5	68030 CPU: 2 68040 CPU: 3 68060 CPU: 5	(msgs/hr * 45 bytes/msg * 8 bits/byte)/(3600 sec/hr)=_____
Meridian Mail voice mail	68030 CPU: 3 68040 CPU: 5 68060 CPU: 8	68030 CPU: 3 68040 CPU: 5 68060 CPU: 8	(msgs/hr * 38.5 bytes/msg * 8 bits/byte)/(3600 sec/hr)=_____
Meridian Mail voice menu	68030 CPU: 5 68040 CPU: 8 68060 CPU: 12	68030 CPU: 5 68040 CPU: 8 68060 CPU: 12	(msgs/hr * 38.5 bytes/msg * 8 bits/byte)/(3600 sec/hr)=_____
Meridian Mail announcements	68030 CPU: 7 68040 CPU: 10 68060 CPU: 15	68030 CPU: 7 68040 CPU: 10 68060 CPU: 15	(msgs/hr * 38.5 bytes/msg * 8 bits/byte)/(3600 sec/hr)=_____

For Meridian Mail 1 through Meridian Mail 9, the CSL link was 4800 baud. Beginning with Meridian Mail 10, the link is 9600 baud. Meridian Mail 11 supports a maximum of 96 ports. Previous releases supported 48 ports.

SDI Applications

In the HSL analysis, include live agents, automated agents, and Meridian Mail agents in the agent total. This will compensate for the assumption of simple calls, since transferred calls will generate additional MAX messages.

Table 81
MSDL real time requirements for SDI applications

SDI	calls/hr A	msgs/call B	msgs/hr C=A*B	msec/msg D	msec E=C*D
CDR	calls/hr with reports= _____	FCDR=old:100 CDR=new: 213	_____	0.05	_____
HSL- Meridian MAX	agents * calls/agent/hr (18.3)= _____	5	_____	8.8	_____
TTY	NA	NA	15000	0.05	_____

There are no traffic reports that provide information on the number of SDI messages directly. For CDR records, determine whether CDR is enabled for incoming, outgoing, and/or internal calls. The number of incoming, outgoing, internal, and tandem calls is available from TFC001. Tandem calls are considered both incoming and outgoing. Alternatively, the number of CDR records can be counted directly. MAX reports can also be counted directly.

Table 82
MSDL peak buffer requirements for SDI applications


SDI	Outgoing	Incoming	Minimum Baud Rate
CDR	30 if baud rate is less than recommended in Table 70 otherwise 68030 CPU: 10 68040 CPU: 15 68060 CPU: 20	1	$(\text{msgs/hr} * 10 \text{ bits/msg}) / (3600 \text{ sec/hr}) =$ _____
HSL - Meridian MAX	Msgs per call simple: 5 medium: 10 complex: 15	1	$(\text{msgs/hr} * 20 \text{ bytes/msg} * 9 \text{ bits/byte}) / (3600 \text{ sec/hr}) =$ _____
TTY	10	10	

Examples

NACD Network with CDR Reports

Consider an NACD network with the topology given in [Figure 11](#). The call flow is provided, where arrows indicate where calls enter the network and where they are answered.

Figure 11
NACD network



Each node has a single ACD DN and calls are queued to the network target DN as soon as they arrive.

For this network, we wish to determine whether a single MSDL on Node B can support DCH1, DCH2, and an SDI port for CDR records on Port 0.

Since we have detailed call flow information, we can develop a messaging model for DCH1 and DCH2 (see Table 83).

Table 83
NACD Message Model

Originating Node	Total Queued	Queued and answered	Queued but not answered	Total messages	DCH1	DCH2
Node A to Node B	3000	500	2500	15500	x	x
Node A to Node C	3000	500	2500	15500	x	x
Node B to Node A	2600	100	2500	11100	x	
Node B to Node C	2600	500	2100	13900		x
Node C to Node A	1650	50	1600	6950	x	x
Node C to Node B	1650	100	1550	7300	x	x

The DCH1 and DCH2 columns indicate whether the messages should be included in the DCH1 and DCH2 message count, respectively. For each row, multiply the entry in the “Queued and answered” column by 11 messages and multiply the entry in the “Queued but not answered” column by 4 messages. The sum of these two values is provided in the “Total messages” column. By summing the rows which should be included for DCH1 and DCH2, we derive the total messages for DCH1: 56350 msg/hr and DCH2: 59150 msg/hr. Note that these messages do not include the impact of CRQS and FCTH which are beyond the scope of this analysis (see Table 77).

Table 84
MSDL real time requirements for D-channel applications

DCH	calls/hr A	msgs/call B	msgs/hr C=A*B	msec/msg D	msec E=C*D
NACD DCH1	NA	NA	56350	pre-R20: 8.8	495880
NACD DCH2	NA	NA	59150	pre-R20: 8.8	520520

Assuming that no non-NACD calls are carried, Node B carries 3750 calls/hour.

Table 85
MSDL real time requirements for SDI applications

SDI	calls/hr A	msgs/call B	msgs/hr C=A*B	msec/ms g D	msec E=C*D
CDR	calls/hr with reports=3750	FCDR=old: 100 FCDR=new: 213	798750 (FCDR=new)	0.05	39938

The total MSDL requirements can then be computed:

Table 86
MSDL Engineering Worksheet

Port	Application	Real Time Required	Peak Buffer Usage Outgoing	Peak Buffer Usage Incoming
0	CDR	39938	10	1
1	DCH-NACD	495880	7	10
2	DCH-NACD	520520	7	10
3				
Total		1056338	24	21

The projected MSDL utilization is $1056338 / 2770000 = 38\%$. Assuming low network congestion, incoming and outgoing peak buffer usage are below 60, so a single MSDL is able to support this configuration. However, due to the potentially high messaging impact of NACD, this MSDL should be re-engineered periodically to determine whether the call volumes or call flow patterns have changed.

MAX, CCR, and D-channel

An Option 81C call center with a 68040 processor which carries 10,000 inbound calls per busy hour would like to configure MAX, AML, a D-channel that provides signaling for 5 T1s, and a D-channel that provides signaling for 3 T1s on a single MSDL. Of the 10,000 inbound calls, 60% receive medium complexity CCR treatment with 40 ACD DN's reporting statistics to the CCR module. Can the configuration be supported?

In the case of the D-channels, assume that a back-up D-channel is configured, so that the number of trunks is $(5 * 24) - 2 = 118$ and $(3 * 24) - 2 = 70$, for the first and second D-channel respectively.

Table 87
MSDL real time requirements

	calls/hr A	msgs/call B	msgs/hr C=A*B	msec/msg D	msec E=C*D
HSL- Meridian MAX	10000	5	50000	8.8	440000
CCR	6000	simple: 10 medium: 20 complex: 30	120000 + (900 * 40) = 156000	7.2	112320 0
ISDN Network	trunks/DCH * calls/hr/trunk (15.6)= 1841	8	14728	pre-R20: 8.8 R20+: 19.2	129606
ISDN Network	trunks/DCH * calls/hr/trunk (15.6)= 1092	8	8736	pre-R20: 8.8 R20+: 19.2	76877

Table 88
MSDL Engineering Worksheet

Port	Application	Real Time Required	Peak Buffer Usage Outgoing	Peak Buffer Usage Incoming
0	Meridian MAX	440000	5	1
1	CCR	1123200	40	10
2	DCH	129606	10	10
3	DCH	76877	10	10
Total		1769683	65	31

While the MSDL supporting this configuration is projected to operate at only $1769683 / 2770000 = 64\%$ of real time capacity, there is a concern that link delays may be a problem due to peakedness of outgoing traffic. It is recommended that the AML or a D-channel be off-loaded to another MSDL.

